

Prior-independent Auctions for Risk-averse Agents

Hu Fu

Jason Hartline

Darrell Hoy

December 11, 2013

We study simple and approximately optimal auctions for agents with a particular form of risk-averse preferences. We show that, for symmetric agents, the optimal revenue (given a prior distribution over the agent preferences) can be approximated by the first-price auction (which is prior independent), and, for asymmetric agents, the optimal revenue can be approximated by an auction with simple form. These results are based on two technical methods. The first is for upper-bounding the revenue from a risk-averse agent. The second gives a payment identity for mechanisms with pay-your-bid semantics.

1 Introduction

We study optimal and approximately optimal auctions for agents with risk-averse preferences. The economics literature on this subject is largely focused on either comparative statics, i.e., is the first-price or second-price auction better when agents are risk averse, or deriving the optimal auction, e.g., using techniques from optimal control, for specific distributions of agent preferences. The former says nothing about optimality but considers realistic prior-independent auctions; the latter says nothing about realistic and prior-independent auctions. Our goal is to study approximately optimal auctions for risk-averse agents that are realistic and not dependent on assumptions on the specific form of the distribution of agent preferences. One of our main conclusions is that, while the second-price auction can be very far from optimal for risk-averse agents, the first-price auction is approximately optimal for an interesting class of risk-averse preferences.

The microeconomic treatment of risk aversion in auction theory suggests that the form of the optimal auction is very dependent on precise modeling details of the preferences of agents, see, e.g., Maskin and Riley (1984) and Matthews (1984). The resulting auctions are unrealistic because of their reliance on the prior assumption and because they are complex (cf. Wilson, 1987). Approximation can address both issues. There may be a class of mechanisms that is simple, natural, and much less dependent on exact properties of the distribution. As an example of this agenda for risk neutral agents, Hartline and Roughgarden (2009) showed that for a large class of distributional assumptions the second-price auction with a reserve is a constant approximation to the optimal single-item auction. This implies that the only information about the distribution of preferences that is necessary for a good approximation is a single number, i.e., a good reserve price. Often from this sort of “simple versus optimal” result it is possible to do away with the reserve price entirely. Dhangwatnotai et al. (2010) and Roughgarden et al. (2012) show that simple and natural mechanisms are approximately optimal quite broadly. We extend this agenda to auction theory for risk-averse agents.

The least controversial approach for modeling risk-averse agent preferences is to assume agents are endowed with a concave function that maps their wealth to a utility. This introduces a non-linearity into the incentive constraints of the agents which in most cases makes auction design analytically intractable. We therefore restrict attention to a very specific form of risk aversion that is both computationally and analytically tractable: utility functions that are linear up to a given capacity and then flat. Importantly, an agent with such a utility function will not trade off a higher probability of winning for a lower price when the utility from such a lower price is greater than her capacity. While capacitated utility functions are unrealistic, they form a basis for general concave utility functions. In our analyses we will endow the benchmark optimal auction with knowledge of the agents’ value distribution and capacity; however, some of the mechanisms we design to approximate this benchmark will be oblivious to them.

As an illustrative example, consider the problem of maximizing welfare by a single-item auction when agents have known capacitated utility functions (but unknown values). Recall that for risk-neutral agents the second-price auction is welfare-optimal as the payments are transfers from the agents to the mechanism and cancel from the objective welfare which is thus equal to value of the winner. (The auctioneer is assumed to have linear utility.) For agents with capacitated utility, the second-price auction can be far from optimal. For instance, when the difference between the highest and second highest bid is much larger than the capacity then the excess value (beyond the capacity) that is received by the winner does not translate to extra utility because it is truncated at the capacity. Instead, a variant of the second-price auction, where the highest bidder wins

and is charged the maximum of the second highest bid and her bid less her capacity, obtains the optimal welfare. Unfortunately, this auction is parameterized by the form of the utility function of the agents. There is, however, an auction, not dependent on specific knowledge of the utility functions or prior distribution, that is also welfare optimal: If the agents values are drawn i.i.d. from a common prior distribution then the first-price auction is welfare-optimal. To see this: (a) standard analyses show that at equilibrium the highest-valued agent wins, and (b) no agent will shade her bid more than her capacity as she receives no increased utility from such a lower payment but her probability of winning strictly decreases.

Our main goal is to duplicate the above observation for the objective of revenue. It is easy to see that the gap between the optimal revenues for risk-neutral and capacitated agents can be of the same order as the gap between the optimal welfare and the optimal revenue (which can be unbounded). When the capacities are small the revenue of the welfare-optimal auction for capacitated utilities is close to its welfare (the winners utility is at most her capacity). Of course, when capacities are infinite or very large then the risk-neutral optimal revenue is close to the capacitated optimal revenue (the capacities are not binding). One of our main technical results shows that even for mid-range capacities one of these two mechanisms that are optimal at the extremes is close to optimal.

As a first step towards understanding profit maximization for capacitated agents, we characterize the optimal auction for agents with capacitated utility functions. We then give a “simple versus optimal” result showing that either the revenue-optimal auction for risk-neutral agents or the above welfare-optimal auction for capacitated agents is a good approximation to the revenue-optimal auction for capacitated agents. The Bulow-Klemperer (1996) Theorem implies that with enough competition (and mild distributional assumptions) welfare-optimal auctions are approximately revenue-optimal. Of course, the first-price auction is welfare-optimal and prior-independent; therefore we conclude that it is approximately revenue-optimal for capacitated agents.

Our “simple versus optimal” result comes from an upper bound on the expected payment of an agent in terms of her allocation rule (cf. Myerson, 1981). This upper bound is the most technical result in the paper; the difficulties that must be overcome by our analysis are exemplified by the following observations. First, unlike in risk-neutral mechanism design, Bayes-Nash equilibrium does not imply monotonicity of allocation rules. There are mechanisms where an agent with a high value would prefer less overall probability of service than she would have obtained if she had a lower value (Example 3.1 in Section 3). Second, even in the case where the capacity is higher than the maximum value of any agent, the optimal mechanism for risk-averse agents can generally obtain more revenue than the optimal mechanism for risk-neutral agents (Example 4.2 in Section 4). This may be surprising because, in such a case, the revenue-optimal mechanism for risk-neutral agents would give any agent a wealth that is within the linear part of her utility function. Finally, while our upper bound on risk-averse payments implies that this relative improvement is bounded by a factor of two for large capacities, it can be arbitrarily large for small capacities (Example 4.1 in Section 4).

It is natural to conjecture that the first-price auction will continue to perform nearly optimally well beyond our simple model (capacitated utility) of risk-averse preferences. It is a relatively straightforward calculation to see that for a large class of risk-averse utility functions from the literature (e.g., Matthews, 1984) the first-price auction is approximately optimal at extremal risk parameters (risk-neutral or extremely risk-averse). We leave to future work the extension of our analysis to mid-range risk parameters for these other families of risk-averse utility functions.

It is significant and deliberate that our main theorem is about the first-price auction which is well known to not have a truthtelling equilibrium. Our goal is a prior-independent mechanism. In particular, we would like our mechanism to be parameterized neither by the distribution on agent preference nor by the capacity that governs the agents utility function. While it is standard in mechanism design and analysis to invoke the *revelation principle* (cf. Myerson, 1981) and restrict attention to auctions with truthtelling as equilibrium, this principle cannot be applied in prior-independent auction design. An auction with good equilibrium can be implemented by one with truthtelling as an equilibrium if the agent strategies can be simulated by the auction. In a Bayesian environment, agent strategies are parameterized by the prior distribution and therefore the suggested revelation mechanism is not generally prior independent.

Risk Aversion, Universal Truthfulness, and Truthfulness in Expectation. Our results have an important implication on a prevailing and questionable perspective that is explicit and implicit broadly in the field of algorithmic mechanism design. Two standard solution concepts from algorithmic mechanism design are “universal truthfulness” and “truthfulness in expectation.” A mechanism is universally truthful if an agent’s optimal (and dominant) strategy is to reveal her values for the various outcomes of the mechanism regardless of the reports of other agents or random coins flipped by the mechanism. In contrast, in a truthful-in-expectation mechanism, revealing truthfully her values only maximizes the agent’s utility in expectation over the random coins tossed by the mechanism. Therefore, a risk-averse agent modeled by a non-linear utility function may not bid truthfully in a truthful-in-expectation mechanism designed for risk-neutral agents, whereas in a universally truthful mechanism an agent behaves the same regardless of her risk attitude. For this reason, the above-mentioned perspective sees universally truthful mechanisms superior because the performance guarantees shown for risk-neutral agents seem to apply to risk-averse agents as well.

This perspective is incorrect because the optimal performance possible by a mechanism is different for risk-neutral and risk-averse agents. In some cases, a mechanism may exploit the risk attitude of the agents to achieve objectives better than the optimal possible for risk-neutral agents; in other cases, the objective itself relies on the utility functions (e.g. social welfare maximization), and therefore the same outcome has a different objective value. In all these situations, the performance guarantee of universally truthful mechanisms measured by the risk-neutral optimality loses its meaning. We have already discussed above two examples for capacitated agents that illustrate this point: for welfare maximization the second-price auction is not optimal, for revenue maximization the risk-neutral revenue-optimal auction can be far from optimal.

The conclusion of the discussion above is that the universally truthful mechanisms from the literature are not generally good when agents are risk averse; therefore, the solution concept of universal truthfulness buys no additional guarantees over truthfulness in expectation. Nonetheless, our results suggest that it may be possible to develop a general theory for prior-independent mechanisms for risk-averse agents. By necessity, though, this theory will look different from the existing theory of algorithmic mechanism design.

Summary of Results. Our main theorem is that the first-price auction is a prior-independent 5-approximation for revenue for two or more agents with i.i.d. values and risk-averse preferences (given by a common capacity). The technical results that enable this theorem are as follows:

- The optimal auction for agents with capacitated utilities is a two-priced mechanism where a

winning agent either pays her full value or her value less her capacity.

- The expected revenue of an agent with capacitated utility and regular value distribution can be bounded in terms of an expected (risk-averse) virtual surplus, where the (risk-averse) virtual value is twice the risk-neutral virtual value plus the value minus capacity (if positive).
- Either the mechanism that optimizes value minus capacity (and charges the Clarke payments or value minus capacity, whichever is higher) or the risk-neutral revenue optimal mechanism is a 3-approximation to the revenue optimal auction for capacitated utilities.
- We characterize the Bayes-Nash equilibria of auctions with capacitated agents where each bidder's payment when served is a deterministic function of her value. An example of this is the first-price auction. The BNE strategies of the capacitated agents can be calculated formulaically from the BNE strategies of risk-neutral agents.

Some of these results extend beyond single-item auctions. In particular, the characterization of equilibrium in the first-price auction holds for position auction environments (i.e., where agents are greedily by bid assigned to positions with decreasing probabilities of service and charged their bid if served). Our simple-versus-optimal 3-approximation holds generally for downward-closed environments, non-identical distributions, and non-identical capacities.

Related Work. The comparative performance of first- and second-price auctions in the presence of risk aversion has been well studied in the Economics literature. From a revenue perspective, first-price auctions are shown to outperform second-price auctions very broadly. Riley and Samuelson (1981) and Holt (1980) show this for symmetric settings where bidders have the same concave utility function. Maskin and Riley (1984) show this for more general preferences.

Matthews (1987) shows that in addition to the revenue dominance, bidders whose risk attitudes exhibit *constant absolute risk aversion (CARA)* are indifferent between first- and second-price auctions, even though they pay more in expectation in the first-price auction. Hu et al. (2010) considers the optimal reserve prices to set in each, and shows that the optimal reserve in the first price auction is less than that in the second price auction. Interestingly, under light conditions on the utility functions, as risk aversion increases, the optimal first-price reserve price decreases.

Matthews (1983) and Maskin and Riley (1984) have considered optimal mechanisms for a single item, with symmetric bidders (i.i.d. values and identical utility function), for CARA and more general preferences.

Recently, Dughmi and Peres (2012) have shown that by insuring bidders against uncertainty, any truthful-in-expectation mechanism for risk-neutral agents can be converted into a dominant-strategy incentive compatible mechanism for risk-averse buyers with no loss of revenue. However, there is potentially much to gain—mechanisms for risk-averse buyers can achieve unboundedly more welfare and revenue than mechanisms for risk-neutral bidders, as we show in Example 4.1 of Section 4.

2 Preliminaries

Risk-averse Agents. Consider selling an item to an agent who has a private valuation v drawn from a known distribution F . Denote the outcome by (x, p) , where $x \in \{0, 1\}$ indicates whether

the agent gets the item, and p is the payment made. The agent obtains a *wealth* of $vx - p$ for such an outcome and the agent's utility is given by a concave utility function $u(\cdot)$ that maps her wealth to utility, i.e., her utility for outcome (x, p) is $u(vx - p)$. Concave utility functions are a standard approach for modeling risk-aversion.¹

A *capacitated* utility function is $u_C(z) = \min(z, C)$ for a given C which we refer to as the *capacity*. Intuitively, small C corresponds to severe risk aversion; large C corresponds to mild risk aversion; and $C = \infty$ corresponds to risk neutrality. An agent views an auction as a deterministic rule that maps a random source and the (possibly random) reports of other agents which we summarize by π , and the report b of the agent, to an allocation and payment. We denote these coupled allocation and payment rules as $x^\pi(b)$ and $p^\pi(b)$, respectively. The agent wishes to maximize her expected utility which is given by $\mathbf{E}_\pi[u_C(vx^\pi(b) - p^\pi(b))]$, i.e., she is a von Neumann-Morgenstern utility maximizer.

Incentives. A strategy profile of agents is $\mathbf{s} = (s_1, \dots, s_n)$ mapping values to reports. Such a strategy profile is in *Bayes-Nash equilibrium* (BNE) if each agent i maximizes her utility by reporting $s_i(v_i)$. I.e., for all i , v_i , and z :

$$\mathbf{E}_\pi [u(v_i x_i^\pi(s_i(v_i)) - p_i^\pi(s_i(v_i)))] \geq \mathbf{E}_\pi [u(v_i x_i^\pi(z) - p_i^\pi(z))]$$

where π denotes the random bits accessed by the mechanism as well as the random inputs $s_j(v_j)$ for $j \neq i$ and $v_j \sim F_j$. A mechanism is *Bayesian incentive compatible* (BIC) if truthtelling is a Bayes-Nash equilibrium: for all i , v_i , and z

$$\mathbf{E}_\pi [u(v_i x_i^\pi(v_i) - p_i^\pi(v_i))] \geq \mathbf{E}_\pi [u(v_i x_i^\pi(z) - p_i^\pi(z))] \quad (\text{IC})$$

where π denotes the random bits accessed by the mechanism as well as the random inputs $v_j \sim F_j$ for $j \neq i$.

We will consider only mechanisms where losers have no payments, and winners pay at most their bids. These constraints imply ex post *individual rationality* (IR). Formulaically, for all i , v_i , and π , $p_i^\pi(v_i) \leq v_i$ when $x_i^\pi(v_i) = 1$ and $p_i^\pi(v_i) = 0$ when $x_i^\pi(v_i) = 0$.

Auctions and Objectives. The revenue of an auction \mathcal{M} is the total payment of all agents; its expected revenue for implicit distribution F and Bayes-Nash equilibrium is denoted $\text{Rev}(\mathcal{M}) = \mathbf{E}_{\pi, \mathbf{v}}[\sum_i p_i^\pi(v_i)]$. The welfare of an auction \mathcal{M} is the total utility of all participants including the auctioneer; its expected welfare is denoted $\text{Welfare}(\mathcal{M}) = \text{Rev}(\mathcal{M}) + \mathbf{E}_{\pi, \mathbf{v}}[\sum_i u(v_i x_i^\pi(v_i) - p_i^\pi(v_i))]$.

Some examples of auctions are: the *first-price auction* (FPA) serves the agent with the highest bid and charges her her bid; the *second-price auction* (SPA) serves the agent with the highest bid and charges her the second-highest bid. The second price auction is incentive compatible regardless of agents' risk attitudes. The *capacitated second-price auction* (CSP) serves the agent with the highest bid and charges her the maximum of her value less her capacity and the second highest bid. The second-price auction for capacitated agents is incentive compatible for capacitated agents because, relative to the second-price auction, the utility an agent receives for truthtelling is unaffected and the utility she receives for any misreport is only (weakly) lower.

¹There are other definitions of risk aversion; this one is the least controversial. See Mas-Colell et al. (1995) for a thorough exposition of expected utility theory.

Two-Priced Auctions. The following class of auctions will be relevant for agents with capacitated utility functions.

Definition 2.1. A mechanism \mathcal{M} is *two-priced* if, whenever \mathcal{M} serves an agent with capacity C and value v , the agent's payment is either v or $v - C$; and otherwise (when not served) her payment is zero. Denote by $x_{\text{val}}(v)$ and $x_C(v)$ probability of paying v and $v - C$, respectively.

Note that from an agent's perspective the outcome of a two-priced mechanism is fully described by a x_C and x_{val} .

Auction Theory for Risk-neutral Agents. For risk neutral agents, i.e., with $u(\cdot)$ equal to the identity function, only the probability of winning and expected payment are relevant. The *interim allocation rule* and *interim payment rule* are given by the expectation of x^π and p^π over π and denoted as $x(b) = \mathbf{E}_\pi[x^\pi(b)]$ and $p(b) = \mathbf{E}_\pi[p^\pi(b)]$, respectively (recall that π encodes the randomization of the mechanism and the reports of other agents).

For risk-neutral agents, Myerson (1981) characterized interim allocation and payment rules that arise in BNE and solved for the revenue optimal auction. These results are summarized in the following theorem.

Theorem 2.1 (Myerson, 1981). *For risk neutral bidders with valuations drawn independently and identically from F ,*

- (a) (*monotonicity*) *The allocation rule $x(v)$ for each agent is monotone non-decreasing in v .*
- (b) (*payment identity*) *The payment rule satisfies $p(v) = vx(v) - \int_0^v x(z)dz$.*
- (c) (*virtual value*) *The ex ante expected payment of an agent is $\mathbf{E}_v[p(v)] = \mathbf{E}_v[\varphi(v)x(v)]$ where $\varphi(v) = v - \frac{1-F(v)}{f(v)}$ is the virtual value for value v .*
- (d) (*optimality*) *When the distribution F is regular, i.e., $\varphi(v)$ is monotone, the second-price auction with reserve $\varphi^{-1}(0)$ is revenue-optimal.*

The payment identity in part (b) implies the *revenue equivalence* between any two auctions with the same BNE allocation rule.

A well-known result by Bulow and Klemperer shows that, in part (d) of Theorem 2.1, instead of having a reserve price to make the second-price auction optimal, one may as well add in another identical bidder to get at least as much revenue.

Theorem 2.2 (Bulow and Klemperer, 1996). *For risk neutral bidders with valuations drawn i.i.d. from a regular distribution, the revenue from the second-price auction with $n + 1$ bidders is at least that of the optimal auction for n bidders.*

3 The Optimal Auctions

In this section we study the form of optimal mechanisms for capacitated agents. In Section 3.1, we show that it is without loss of generality to consider two-priced auctions, and in Section 3.2 we characterize the incentive constraints of two-priced auctions. In Section 3.3 we use this characterization to show that the optimal auction (in discrete type spaces) can be computed in polynomial time in the number of types.

3.1 Two-priced Auctions Are Optimal

Recall a two-priced auction is one where when any agent is served she is either charged her value or her value minus her capacity. We show below that restricting our attention to two-priced auctions is without loss for the objective of revenue.

Theorem 3.1. *For any auction on capacitated agents there is a two-priced auction with no lower revenue.*

Proof. We prove this theorem in two steps. In the first step we show, quite simply, that if an agent with a particular value received more wealth than C then we can truncate her wealth to C (by charging her more). With her given value she is indifferent to this change, and for all other values this change makes misreporting this value (weakly) less desirable. Therefore, such a change would not induce misreporting and only (weakly) increases revenue. This first step gives a mechanism wherein every agent's wealth is in the linear part of her utility function. The second step is to show that we can transform the distribution of wealth into a two point distribution. Whenever an agent with value v is offered a price that results in a wealth $w \in [0, C]$, we instead offer her a price of $v - C$ with probability w/C , and a price of v with the remaining probability. Both the expected revenue and the utility of a truthful bidder is unchanged. The expected utility of other types to misreport v , however, weakly decreases by the concavity of u_C , because mixing over endpoints of an interval on a concave function gives less value than mixing over internal points with the same expectation. \square

3.2 Characterization of Two-Priced Auctions

In this section we characterize the incentive constraints of two-priced auctions. We focus on the induced two-priced mechanism for a single agent given the randomization π of other agent values and the mechanism. The interim two-priced allocation rule of this agent is denoted by $x(v) = x_{\text{val}}(v) + x_C(v)$.

Lemma 3.2. *A mechanism with two-price allocation rule $x = x_{\text{val}} + x_C$ is BIC if and only if for all v and v^+ such that $v < v^+ \leq v + C$,*

$$\frac{x_{\text{val}}(v)}{C} \leq \frac{x_C(v^+) - x_C(v)}{v^+ - v} \leq \frac{x(v^+)}{C}. \quad (1)$$

Equation (1) can be equivalently written as the following two linear constraints on x_C , for all $v^- \leq v \leq v^+ \in [v - C, v + C]$:

$$x_C(v^+) \geq x_C(v) + \frac{v^+ - v}{C} \cdot x_{\text{val}}(v), \quad (2)$$

$$x_C(v^-) \geq x_C(v) - \frac{v - v^-}{C} \cdot x(v). \quad (3)$$

Equations (2) and (3) are illustrated in Figure 1. For a fixed v , (2) with $v^+ = v + C$ yields a lower bounding line segment from $(v, x_C(v))$ to $(v + C, x_C(v) + x_{\text{val}}(v))$, and (3) with $v^- = v - C$ gives a lower bounding line segment from $(v, x_C(v))$ to $(v - C, x_C(v) - x(v))$. Note that (2) implies that x_C is monotone.

In the special case when x_C is differentiable, by taking v^+ approaching v in (1), we have $\frac{x_{\text{val}}(v)}{C} \leq x'_C(v) \leq \frac{x(v)}{C}$ for all v . In general, we have the following condition in the integral form (see Appendix A for a proof).

Corollary 3.3. *The allocation rule $x = x_{\text{val}} + x_C$ of a BIC two-priced mechanism for all $v < v^+$ satisfies:*

$$\int_v^{v^+} \frac{x_{\text{val}}(z)}{C} dz \leq x_C(v^+) - x_C(v) \leq \int_v^{v^+} \frac{x(z)}{C} dz. \quad (4)$$




Figure 1: Fixing $x(v) = x_{\text{val}}(v) + x_C(v)$, the dashed line between points $(v - C, x_C(v) - x(v))$, $(v, x_C(v))$, and $(v + C, x_C(v) + x_{\text{val}}(v))$ (denoted by “•”) depicts the lower bounds from (2) and (3) on x_C for values in $[v - C, v + C]$.

Importantly, the equilibrium characterization of two-priced mechanisms does not imply monotonicity of the allocation rule x . This is in contrast with mechanisms for risk-neutral agents, where incentive compatibility requires a monotone allocation rule (Theorem 2.1, part (a)). This non-monotonicity is exhibited in the following example.

Example 3.1. There is a single-agent two-priced mechanism with a non-monotone allocation rule. Our agent has two possible values $v = 3$ and $v = 4$, and capacity C of 2. We give a two price mechanism. Recall that $x_C(v)$ is the probability with which the mechanism sells the item and charges $v - C$; $x_{\text{val}}(v)$ is the probability with which the mechanism sells the item and charges v ; and $x(v) = x_C(v) + x_{\text{val}}(v)$. The mechanism and its outcome are summarized in the following table.

v	x	x_C	x_{val}	utility from truthful reporting	utility from misreporting
3	5/6	1/2	1/3	1	2/3
4	2/3	2/3	0	4/3	4/3

3.3 Optimal Auction Computation

Solving for the optimal mechanism is computationally tractable for any discrete (explicitly given) type space T . Given a discrete valuation distribution on support T , one can use $2|T|$ variables to represent the allocation rule of any two-priced mechanism, and the expected revenue is a linear sum of these variables. Lemma 3.2 shows that one can use $O(|T|^2)$ linear constraints to express all BIC allocations, and hence the revenue optimization for a single bidder can be solved by a $O(|T|^2)$ -sized linear program. Furthermore, using techniques developed by Cai et al. (2012) and

Alaei et al. (2012), in particular the “token-passing” characterization of single-item auctions by Alaei et al. (2012), we obtain:

Theorem 3.4. *For n bidders with independent valuations with type spaces T_1, \dots, T_n and capacities C_1, \dots, C_n , one can solve for the optimal single-item auction with a linear program of size $O\left(\left(\sum_i |T_i|\right)^2\right)$.*

4 An Upper Bound on Two-Priced Expected Payment

In this section we will prove an upper bound on the expected payment from any capacitated agent in a two-priced mechanism. This upper bound is analogous in purpose to the identity between expected risk-neutral payments and expected virtual surplus of Myerson (1981) from which optimal auctions for risk-neutral agents are derived. We use this bound in Section 5.2 and Section 5.3 to derive approximately optimal mechanisms.

As before, we focus on the induced two-priced mechanism for a single agent given the randomization π of other agent values and the mechanism. The expected payment of a bidder of value v under allocation rule $x(v) = x_C(v) + x_{\text{val}}(v)$ is $p(v) = v \cdot x_{\text{val}}(v) + (v - C) \cdot x_C(v) = v \cdot x(v) - C \cdot x_C(v)$.

Recall from Theorem 2.1 that the (risk-neutral) virtual value for an agent with value drawn from distribution F is $\varphi(v) = v - \frac{1-F(v)}{f(v)}$ and that the expected risk-neutral payment for allocation rule $x(\cdot)$ is $\mathbf{E}_v[\varphi(v)x(v)]$. Denote $\max(0, \varphi(v))$ by $\varphi^+(v)$ and $\max(v - C, 0)$ by $(v - C)^+$.

Theorem 4.1. *For any agent with value $v \sim F$, capacity C , and two-priced allocation rule $x(v) = x_{\text{val}}(v) + x_C(v)$,*

$$\begin{aligned} \mathbf{E}_v [p(v)] &\leq \mathbf{E}_v \left[\varphi^+(v) \cdot x(v) \right] + \mathbf{E}_v \left[\varphi^+(v) \cdot x_C(v) \right] \\ &\quad + \mathbf{E}_v \left[(v - C)^+ \cdot x_C(v) \right]. \end{aligned}$$

Corollary 4.2. *When bidders have regular distributions and a common capacity, either the risk-neutral optimal auction or the capacitated second price auction (whichever has higher revenue) gives a 3-approximation to the optimal revenue for capacitated agents.*

Proof. For each of the three parts of the revenue upper bound of Theorem 4.1, there is a simple auction that optimizes the expectation of the part across all agents. For the first two parts, the allocation rules across agents (both for $x(\cdot)$ and $x_C(\cdot)$) are feasible. When the distributions of agent values are regular (i.e., the virtual value functions are monotone), the risk-neutral revenue-optimal auction optimizes virtual surplus across all feasible allocations (i.e., expected virtual value of the agent served); therefore, its expected revenue upper bounds the first and second parts of the bound in Theorem 4.1. The revenue of the third part is again the expectation of a monotone function (in this case $(v - C)^+$) times the service probability. The auction that serves the agent with the highest (positive) “value minus capacity” (and charges the winner the maximum of her “minimum winning bid,” i.e., the second-price payment rule, and her “value minus capacity”) optimizes such an expression over all feasible allocations; therefore, its revenue upper bounds this third part of the bound in Theorem 4.1. When capacities are identical, this auction is the capacitated second price auction. \square

Before proving Theorem 4.1, we give two examples. The first shows that the gap between the revenue of the capacitated second-price auction and the risk-neutral revenue-optimal auction (i.e., the two auctions from Corollary 4.2) can be arbitrarily large. This means that there is no hope that an auction for risk-neutral agents always obtains a good revenue for risk-averse agents. The second example shows that even when all values are bounded from above by the capacity (and therefore, capacities are never binding in a risk-neutral auction) an auction for risk-averse agents can still take advantage of risk aversion to generate higher revenue. Consequently, the fact that we have two risk-neutral revenue terms in the bound of Theorem 4.1 is necessary (as the “value minus capacity” term is zero in this case).

Example 4.1. The *equal revenue* distribution on interval $[1, h]$ has distribution function $F(z) = 1 - 1/z$ (with a point mass at h). The distribution gets its name because such an agent would accept any offer price of p with probability $1/p$ and generate an expected revenue of one. With one such agent the optimal risk-neutral revenue is one. Of course, an agent with capacity $C = 1$ would happily pay her value minus her capacity to win all the time (i.e., $x(v) = x_C(v) = 1$). The revenue of this auction is $\mathbf{E}[v] - 1 = \ln h$. For large h , this is unboundedly larger than the revenue we can obtain from a risk-neutral agent with the same distribution.

Example 4.2. The revenue from a two-priced mechanism can be better than the optimal risk-neutral revenue even when all values are no more than the capacity. Consider selling to an agent with capacity of $C = 1000$ and value drawn from the equal revenue distribution from Example 4.1 with $h = 1000$.

The following two-priced rule is BIC and generates revenue of approximately 1.55 when selling to such a bidder. Let $x_C(v) = \frac{0.6}{1000}(v - 1)$, $x(v) = \min(x_C(v) + 0.6, 1)$, and $x_{\text{val}}(v) = x(v) - x_C(v)$ (shown in Figure 2). Recall that the expected payment from an agent with value v can be written as $vx(v) - Cx_C(v)$; for small values, this will be approximately 0.6; for large values this will increase to 400. The expected revenue is $\int_1^{1000} (z \cdot x(z) - 1000 x_C(z)) f(z) dz + \frac{1}{1000}(1000 \cdot x_{\text{val}}(1000)) \approx 1.15 + 0.4 \approx 1.55$, an improvement over the optimal risk-neutral revenue of 1.




Figure 2: With $C = 1000$ and values from the equal revenue distribution on $[1, 1000]$, this two-priced mechanism is BIC and achieves 1.55 times the revenue of the optimal risk-neutral mechanism.

In the remainder of this section we instantiate the following outline for the proof of Theorem 4.1. First, we transform any given two-priced allocation rule $x = x_{\text{val}} + x_C$ into a new two-priced rule $\bar{x}(v) = \bar{x}_C(v) + \bar{x}_{\text{val}}(v)$ (for which the expected payment is $\bar{p}(v) = v\bar{x}(v) - C\bar{x}_C(v)$). While this

transformation may violate some incentive constraints (from Lemma 3.2), it enforces convexity of $\bar{x}_C(v)$ on $v \in [0, C]$ and (weakly) improves revenue. Second, we derive a simple upper bound on the payment rule $\bar{p}(\cdot)$. Finally, we use the enforced convexity property of $\bar{x}_C(\cdot)$ and the revenue upper bound to partition the expected payment $\mathbf{E}_v[\bar{p}(v)]$ by the three terms that can each be attained by simple mechanisms.

4.1 Two-Priced Allocation Construction

We now construct a two-priced allocation rule $\bar{x} = \bar{x}_{\text{val}} + \bar{x}_C$ from $x = x_{\text{val}} + x_C$ for which (a) revenue is improved, i.e., $\bar{p}(v) \geq p(v)$, and (b) the probability the agent pays her value minus capacity, $\bar{x}_C(v)$, is convex for $v \in [0, C]$. In fact, given x_{val} , \bar{x}_C is the smallest function for which IC constraint (2) holds; and in the special case when x_{val} is monotone, the left-hand side of (4) is tight for \bar{x}_C on $[0, C]$. Other incentive constraints may be violated by \bar{x} , but we use it only as an upper bound for revenue.

Definition 4.1 (\bar{x}). We define $\bar{x} = \bar{x}_C + \bar{x}_{\text{val}}$ as follows:

- (a) $\bar{x}_{\text{val}}(v) = x_{\text{val}}(v)$;
- (b) Let $r(v)$ be $\frac{1}{C} \sup_{z \leq v} x_{\text{val}}(z)$, and let

$$\bar{x}_C(v) = \begin{cases} \int_0^v r(y) \, dy, & v \in [0, C]; \\ x_C(v), & v > C. \end{cases} \quad (5)$$

Lemma 4.3 (Properties of \bar{x}).

- (a) On $v \in [0, C]$, $\bar{x}_C(\cdot)$ is a convex, monotone increasing function.
- (b) On all v , $\bar{x}_C(v) \leq x_C(v)$.
- (c) The incentive constraint from the left-hand side of (4) holds for \bar{x}_C : $\frac{1}{C} \int_v^{v^+} \bar{x}_{\text{val}}(z) \, dz \leq \bar{x}_C(v^+) - \bar{x}_C(v)$ for all $v < v^+$.
- (d) On all v , $\bar{x}_C(v) \leq x_C(v)$, $\bar{x}(v) \leq x(v)$, and $\bar{p}(v) \geq p(v)$.

The proof of part (b) is technical, and we give a sketch here. Recall that, for each v , the IC constraint (2) gives a linear constraint lower bounding $x_C(v^+)$ for every $v^+ > v$. If one decreases $x_C(v)$, the lower bound it imposes on $x_C(v^+)$ is simply “pulled down” and is less binding. The definition of \bar{x}_C simply lands $\bar{x}_C(v)$ on the most binding lower bound, and therefore not only makes $\bar{x}_C(v)$ at most $x_C(v)$, but also lowers the linear constraint that v imposes on larger values. If the number of values is countable or if x_{val} is piecewise constant, the lemma is easy to see by induction. A full proof for the general case of part (b), along with the proofs of the other more direct parts of Lemma 4.3, is given in Appendix B.

4.2 Payment Upper Bound

Recall that $\bar{p}(v)$ is the expected payment corresponding with two-priced allocation rule $\bar{x}(v)$. We now give an upper bound on $\bar{p}(v)$.

Lemma 4.4. *The payment $\bar{p}(v)$ for v and two-priced rule $\bar{x}(v)$ satisfies*

$$\bar{p}(v) \leq v\bar{x}(v) - \int_0^v \bar{x}(z) dz + \int_0^v \bar{x}_C(z) dz. \quad (6)$$


Proof. View a two-priced mechanism $\bar{x} = \bar{x}_{\text{val}} + \bar{x}_C$ as charging v with probability $\bar{x}(v)$ and giving a rebate of C with probability $\bar{x}_C(v)$. We bound this rebate as follows (which proves the lemma):

$$\begin{aligned} C \cdot \bar{x}_C(v) &\geq C \cdot \bar{x}_C(0) + \int_0^v \bar{x}_{\text{val}}(z) dz \\ &\geq \int_0^v \bar{x}(z) dz - \int_0^v \bar{x}_C(z) dz. \end{aligned}$$

The first inequality is from part (c) of Lemma 4.3. The second inequality is from the definition of $\bar{x}_C(0) = 0$ in (5) and $\bar{x}_{\text{val}}(v) = \bar{x}(v) - \bar{x}_C(v)$. See Figure 3 for an illustration. \square

4.3 Three-part Payment Decomposition

Below, we bound $\bar{p}(\cdot)$ (and hence $p(\cdot)$) in terms of the expected payment of three natural mechanisms. As seen geometrically in Figure 4, the bound given in Lemma 4.4 can be broken into two parts: the area above $\bar{x}(\cdot)$, and the area below $\bar{x}_C(\cdot)$. We refer to the former as $\bar{p}^I(\cdot)$; we further split the latter quantity into two parts: $\bar{p}^{II}(\cdot)$, the area corresponding to $v \in [0, C]$, and $\bar{p}^{III}(\cdot)$, that corresponding to $v \in [C, v]$. We define these quantities formally below:



(a) Shaded region is the expected payment from an agent of value v .

(b) Shaded region upper bounds expected payment from an agent with value v , shown in Lemma 4.4.

Figure 3

$$\bar{p}^I(v) = \bar{x}(v)v - \int_0^v \bar{x}(z) dz, \quad (7)$$

$$\bar{p}^{II}(v) = \int_0^{\min\{v, C\}} \bar{x}_C(z) dz \quad (8)$$

$$\bar{p}^{III}(v) = \begin{cases} 0, & v \leq C; \\ \int_C^v \bar{x}_C(z) dz, & v > C. \end{cases} \quad (9)$$

Proof of Theorem 4.1. We now bound the revenue from each of the three parts of the payment decomposition. These bounds, combined with part (d) of Lemma 4.3 and Lemma 4.4, immediately give Theorem 4.1.

Part 1. $\mathbf{E}_v[\bar{p}^I(v)] = \mathbf{E}_v[\varphi(v) \cdot \bar{x}(v)] \leq \mathbf{E}_v[\varphi^+(v) \cdot x(v)]$.

Formulaically, $\bar{p}^I(\cdot)$ corresponds to the risk-neutral payment identity for $\bar{x}(\cdot)$ as specified by part (b) of Theorem 2.1; by part (c) of Theorem 2.1, in expectation over v , this payment is equal to the expected virtual surplus $\mathbf{E}_v[\varphi(v) \cdot \bar{x}(v)]$.² The inequality follows as terms $\varphi(v)$ and $\bar{x}(v)$ in this expectation are point-wise upper bounded by $\varphi^+(v) = \max(\varphi(v), 0)$ and $x(v)$, respectively, the latter by part (d) of Lemma 4.3.

Part 2. $\mathbf{E}_v[\bar{p}^{II}(v)] \leq \mathbf{E}_v[\varphi(v) \cdot \bar{x}_C(v)] \leq \mathbf{E}_v[\varphi^+(v) \cdot x_C(v)]$.

By definition of $\bar{p}^{II}(\cdot)$ in (8), if the statement of the lemma holds for $v = C$ it holds for $v > C$; so we argue it only for $v \in [0, C]$. Formulaically, with respect to a risk-neutral agent with allocation rule $\bar{x}_C(\cdot)$, the risk-neutral payment is $v \cdot \bar{x}_C(v) - \int_0^v \bar{x}_C(z) dz$, the surplus is $v \cdot \bar{x}_C(v)$, and the risk-neutral agent's utility (the difference between the surplus and payment) is $\int_0^v \bar{x}_C(z) dz = \bar{p}^{II}(v)$. Convexity of $\bar{x}_C(\cdot)$, from part (a) of Lemma 4.3, implies that the risk-neutral payment is at least half the surplus, and so is at least the risk-neutral utility. The lemma follows, then, by the same argument as in the previous part.

²Note: This equality does not require monotonicity of the allocation rule $\bar{x}(\cdot)$; as long as part (b) of Theorem 2.1 formulaically holds, part (c) follows from integration by parts.




Figure 4: Breakdown of the expected payment upper bound in a two-priced auction.

Part 3. $\mathbf{E}_v[\bar{p}^{\text{III}}(v)] \leq \mathbf{E}_v[(v - C)^+ \cdot \bar{x}_C(v)] = \mathbf{E}_v[(v - C)^+ \cdot x_C(v)]$.

The statement is trivial for $v \leq C$ so assume $v \geq C$. By definition $\bar{x}_C(v) = x_C(v)$ for $v > C$. By (2), $x_C(\cdot)$ is monotone non-decreasing. Hence, for $v > C$, $\bar{p}^{\text{III}}(v) = \int_C^v x_C(z) dz \leq \int_C^v x_C(v) dz = (v - C) \cdot x_C(v)$. Plugging in $(v - C)^+ = \max(v - C, 0)$ and taking expectation over v , we obtain the bound. □

5 Approximation Mechanisms and a Payment Identity

In this section we first give a payment identity for Bayes-Nash equilibria in mechanisms that charge agents a deterministic amount upon winning (and zero upon losing). Such one-priced payment schemes are not optimal for capacitated agents; however, we will show that they are approximately optimal. When agents are symmetric (with identical distribution and capacity) we use this payment identity to prove that the first-price auction is approximately optimal. When agents are asymmetric we give a simple direct-revelation one-priced mechanism that is BIC and approximately optimal.

5.1 A One-price Payment Identity

For risk-neutral agents, the Bayes-Nash equilibrium conditions entail a payment identity: given an interim allocation rule, the payment rule is fixed (Theorem 2.1, part (b)). For risk-averse agents there is no such payment identity: there are mechanisms with the identical BNE allocation rules but distinct BNE payment rules. We restrict attention to auctions wherein an agent's payment is a deterministic function of her value (if she wins) and zero if she loses. We call these *one-priced* mechanisms; for these mechanisms there is a (partial) payment identity.

Payment identities are an interim phenomenon. We consider a single agent and the induced allocation rule she faces from a Bayesian incentive compatible auction (or, by the revelation principle, any BNE of any mechanism). This allocation rule internalizes randomization in the environment and the auction, and specifies the agents' probability of winning, $x(v)$, as a function of her value. Given allocation rule $x(v)$, the risk-neutral expected payment is $p^{\text{RN}}(v) = v \cdot x(v) - \int_0^v x(z) dz$ (Theorem 2.1, part (b)). Given an allocation rule $x(v)$, a one-priced mechanism with payment rule $p(v)$ would charge the agent $p(v)/x(v)$ upon winning and zero otherwise (for an expected payment of $p(v)$). Define $p^{\text{VC}}(v) = (v - C) \cdot x(v)$ which, intuitively, gives a lower bound on a capacitated agent's willingness to trade-off decreased probability of winning for a cheaper price.

Theorem 5.1. *An allocation rule x and payment rule p are the BNE of a one-priced mechanism if and only if (a) x is monotone non-decreasing and (b) if $p(v) \geq p^{\text{VC}}(v)$ for all v then $p = p^C$ is defined as*

$$p^C(0) = 0, \tag{10}$$

$$p^C(v) = \max \left(p^{\text{VC}}(v), \sup_{v^- < v} \left\{ p^C(v^-) + (p^{\text{RN}}(v) - p^{\text{RN}}(v^-)) \right\} \right). \tag{11}$$

Moreover, if x is strictly increasing then $p(v) \geq p^{\text{VC}}(v)$ for all v and $p = p^C$ is the unique equilibrium payment rule.

The payment rule should be thought of in terms of two “regimes”: when $p^C = p^{VC}$, and when $p^C > p^{VC}$, corresponding to the first and second terms in the max argument of (11) respectively. In the latter regime, (11) necessitates that $\frac{d}{dv}p^C(v) = \frac{d}{dv}p^{RN}(v)$; for nearby such points v and $v + \epsilon$, the v^- involved in the supremum will be the same, and thus $p^C(v + \epsilon) - p^C(v) = p^{RN}(v + \epsilon) - p^{RN}(v)$.

The proof is relegated to Appendix C. The main intuition for this characterization is that risk-neutral payments are “memoryless” in the following sense. Suppose we fix $p^{RN}(v)$ for a v and ignore the incentive of an agent with value $v^+ > v$ to prefer reporting $v^- < v$, then the risk-neutral payment for all $v^+ > v$ is $p^{RN}(v^+) = p(v) + \int_v^{v^+} (x(v^+) - x(z)) dz$. This memorylessness is simply the manifestation of the fact that the risk-neutral payment identity imposes local constraints on the derivatives of the payment, i.e., $\frac{d}{dv}p^{RN}(v) = v \cdot \frac{d}{dv}x(v)$.

There is a simple algorithm for constructing the risk-averse payment rule p^C from the risk-neutral payment rule p^{RN} (for the same allocation rule x).

0. For $v < C$, $p^C(v) = p^{RN}(v)$.
- (a) The $p^C(v) = p^{RN}(v)$ identity continues until the value v' where $p^C(v') = p^{VC}(v')$, and $p^C(v)$ switches to follow $p^{VC}(v)$.
- (b) When v increases to the value v'' where $\frac{d}{dv}p^{RN}(v'') = \frac{dp^{VC}}{dv}(v'')$ then $p^C(v)$ switches to follow $p^{RN}(v)$ shifted up by the difference $p^{VC}(v'') - p^{RN}(v'')$ (i.e., its derivative $\frac{d}{dv}p^C(v)$ follows $\frac{d}{dv}p^{RN}(v)$).
- (c) Repeat this process from Step (a).

Lemma 5.2. *The one-priced BIC allocation rule x and payment rule p^C satisfy the following*

- (a) *For all v , $p^C(v) \geq \max(p^{RN}(v), p^{VC}(v))$.*
- (b) *Both $p^C(v)$ and $p^C(v)/x(v)$ are monotone non-decreasing.*

The proof of part (b) is contained in the proof of Lemma C.3 in Appendix C, and part (a) follows directly from equations (10) and (11).

5.2 Approximate Optimality of First-price Auction

We show herein that for agents with a common capacity and values drawn i.i.d. from a continuous, regular distribution F with strictly positive density the first-price auction is approximately optimal.

It is easy to solve for a symmetric equilibrium in the first-price auction with identical agents. First, guess that in BNE the agent with the highest value wins. When the agents are i.i.d. draws from distribution F , the implied allocation rule is $x(v) = F^{n-1}(v)$. Theorem 5.1 then gives the necessary equilibrium payment rule $p^C(v)$ from which the bid function $b^C(v) = p^C(v)/x(v)$ can be calculated. We verify that the initial guess is correct as Lemma 5.2 implies that the bid function is symmetric and monotone. There is no other symmetric equilibrium.³

³Any other symmetric equilibrium must have an allocation rule that is increasing but not always strictly so. For this to occur the bid function must not be strictly increasing implying a point mass in the distribution of bids. Of course, a point mass in a symmetric equilibrium bid function implies that a tie is not a measure zero event. Any agent has a best response to such an equilibrium of bidding just higher than this pointmass so at essentially the same payment, she always “wins” the tie.

Proposition 5.3. *The first-price auction for identical (capacity and value distribution) agents has a unique symmetric BNE wherein the highest valued agent wins.*

The expected revenue at this equilibrium is $n \mathbf{E}_v[p^C(v)]$. Lemma 5.2 implies that p^C is at least p^{RN} and p^{VC} .

Corollary 5.4. *The expected revenue of the first-price auction for identical (capacity and value distribution) agents is at least that of the capacitated second-price auction and at least that of the second-price auction.*

Our main theorem then follows by combining Corollary 5.4 with the revenue bound in Theorem 4.1 and Theorem 2.2 by Bulow and Klemperer (1996).

Theorem 5.5. *For $n \geq 2$ agents with common capacity and values drawn i.i.d. from a regular distribution, the revenue in the first price auction (FPA) in the symmetric Bayes-Nash equilibrium is a 5-approximation to the optimal revenue.*

Proof. An immediate consequence of Theorem 2.2 is that for $n \geq 2$ risk-neutral, regular, i.i.d. bidders, the second-price auction extracts a revenue that is at least half the optimal revenue; hence, by Corollary 4.2, the optimal revenue for capacitated bidders by any BIC mechanism is at most four times the second-price revenue plus the capacitated second-price revenue. Since the first-price auction revenue in BNE for capacitated agents is at least the capacitated second-price revenue and the second-price revenue, the first-price revenue is a 5-approximation to the optimal revenue.⁴ \square

5.3 Approximate Optimality of One-Price Auctions

We now consider the case of asymmetric value distributions and capacities. In such settings the highest-bid-wins first-price auction does not have a symmetric equilibria and arguing revenue bounds for it is difficult. Nonetheless, we can give asymmetric one-priced revelation mechanisms that are BIC and approximately optimal. With respect to Example 4.2 in Section 4 which shows that the option to charge two possible prices from a given type may be necessary for optimal revenue extraction, this result shows that charging two prices over charging one price does not confer a significant advantage.

Theorem 5.6. *For n (non-identical) agents, their capacities C_1, \dots, C_n , and regular value distributions F_1, \dots, F_n , there is a one-priced BIC mechanism whose revenue is at least one third of the optimal (two-priced) revenue.*

Proof. Recall from Theorem 4.1 that either the risk-neutral optimal revenue or $\mathbf{E}_{v_1, \dots, v_n}[\max\{(v_i - C_i)_+\}]$ is at least one third of the optimal revenue. We apply Theorem 5.1 to two monotone allocation rules:

- (a) the interim allocation rule of the risk-neutral optimal auction, and
- (b) the interim allocation rule specified by: serve agent i that maximizes $v_i - C_i$, if positive; otherwise, serve nobody.

⁴In fact, the Bulow and Klemperer (1996) result shows that the second-price auction is asymptotically optimal so for large n this bound can be asymptotically improved to three.

As both allocations are monotone, we apply Theorem 5.1 to obtain two single-priced BIC mechanisms. By Lemma 5.2, the expected revenue of the first mechanism is at least the risk neutral optimal revenue, and the expected revenue of the second mechanism is at least $\mathbf{E}_{v_1, \dots, v_n}[\max\{(v_i - C_i)_+\}]$. The theorem immediately follows for the auction with the higher expected revenue. \square

Although Theorem 5.6 is stated as an existential result, the two one-priced mechanisms in the proof can be described analytically using the algorithm following Theorem 5.1 for calculating the capacitated BIC payment rule. The interim allocation rules are straightforward (the first: $x_i(v_i) = \prod_{j \neq i} F_j(\varphi_j^{-1}(\varphi_i(v_i)))$, and the second: $x_i(v_i) = \prod_{j \neq i} F_j(v_i - C_i + C_j)$), and from these we can solve for $p_i^{C_i}(v)$.

6 Conclusions

For the purpose of keeping the exposition simple, we have applied our analysis only to single-item auctions. Our techniques, however, as they focus on analyzing and bounding revenue of a single agent for a given allocation rule, generalize easily to structurally rich environments. Notice that the main theorems of Sections 3, 4, and the first part of Section 5 do not rely on any assumptions on the feasibility constraint except for downward closure, i.e., that it is feasible to change an allocation by withholding service to an agent who was formerly being served.

For example, our prior-independent 5-approximation result generalizes to symmetric feasibility constraints such as position auctions. A position auction environment is given by a decreasing sequence of weights $\alpha_1, \dots, \alpha_n$ and the first-price position auction assigns the agents to these positions greedily by bid. With probability α_i the agent in position i receives an item and is charged her bid; otherwise she is not charged. (These position auctions have been used to model pay-per-click auctions for selling advertisements on search engines where α_i is the probability that an advertiser gets clicked on when her ad is shown in the i th position on the search results page.) For agents with identical capacities and value distributions, the first-price position auction where the bottom half of the agents are always rejected is a 5-approximation to the revenue-optimal position auction (that may potentially match all the agents to slots).

Our one- versus two-price result generalizes to asymmetric capacities, asymmetric distributions, and asymmetric downward-closed feasibility constraints. A downward-closed feasibility constraint is given by a set system which specifies which subset of agents can be simultaneously served. Downward-closure requires that any subset of a feasible set is feasible. A simple one-priced mechanism is a 3-approximation to the optimal mechanism in such an environment. The mechanism is whichever has higher revenue of the standard (risk neutral) revenue-optimal mechanism (which serves the subset of agents with the highest virtual surplus, i.e., sum of virtual values) and the one-priced revelation mechanism that serves the set of agents S that maximizes $\sum_{i \in S} (v_i - C_i)^+$ subject to feasibility.

A main direction for future work is to relax some of the assumptions of our model. Our approach to optimizing over mechanisms for risk-averse agents relies on (a) the simple model of risk aversion given by capacitated utilities and (b) that losers neither make (i.e., ex post individual rationality) nor receive payments (i.e., no bribes). These restrictions are fundamental for obtaining linear incentive compatibility constraints. Of great interest in future study is relaxation of these assumptions.

There is a relatively well-behaved class of risk attitudes known as *constant absolute risk aversion*

where the utility function is parameterized by risk parameter R as $u_R(w) = \frac{1}{R}(1 - e^{-Rw})$. These model the setting in which a bidder’s risk aversion is independent of wealth, and hence bidders view a lottery over payments for an item the same no matter their valuations. Matthews (1984) exploits this and derives the optimal auction for such risk attitudes. A first step in extending our results to more interesting risk attitudes would be to consider such risk preferences.

Our analytical (and computational) solution to the optimal auction problem for agents with capacitated utilities requires an *ex post individual rationality* constraint on the mechanism that is standard in algorithmic mechanism design. This constraint requires that an agent who loses the auction cannot be charged. While such a constraint is natural in many settings, it is with loss and, in fact, ill motivated for settings with risk-averse agents. One of the most standard mechanisms for agents with risk-averse preferences is the “insurance mechanism” where an agent who may face some large liability with small probability will prefer to pay a constant insurance fee so that the insurance agency will cover the large liability in the event that it is incurred. This mechanism is not *ex post* individually rational. Does the first-price auction (which is *ex post* individual rational) approximate the optimal interim individually rational mechanism?

References

- Alaei, S., Fu, H., Haghpahan, N., Hartline, J., and Malekian, A. (2012). Bayesian optimal auctions via multi- to single-agent reduction. In *ACM Conference on Electronic Commerce*.
- Bulow, J. and Klemperer, P. (1996). Auctions Versus Negotiations. *The American Economic Review*, 86(1):180–194.
- Cai, Y., Daskalakis, C., and Weinberg, S. M. (2012). An algorithmic characterization of multi-dimensional mechanisms. In *STOC*, pages 459–478.
- Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2010). Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138.
- Dughmi, S. and Peres, Y. (2012). Mechanisms for Risk Averse Agents, Without Loss. *arXiv preprint arXiv:1206.2957*, pages 1–9.
- Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234.
- Holt, C. A. (1980). Competitive bidding for contracts under alternative auction procedures. *The Journal of Political Economy*, 88(3):433–445.
- Hu, A., Matthews, S. A., and Zou, L. (2010). Risk aversion and optimal reserve prices in first- and second-price auctions. *Journal of Economic Theory*, 145(3):1188–1202.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.
- Maskin, E. and Riley, J. (1984). Optimal auctions with risk averse buyers. *Econometrica*, 52(6):1473–1518.

- Matthews, S. (1983). Selling to risk averse buyers with unobservable tastes. *Journal of Economic Theory*, 30(2):370–400.
- Matthews, S. A. (1984). On the implementability of reduced form auctions. *Econometrica*, 52(6):pp. 1519–1522.
- Matthews, S. A. (1987). Comparing auctions for risk averse buyers: A buyer’s point of view. *Econometrica*, 55(3):633–646.
- Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):pp. 58–73.
- Riley, J. and Samuelson, W. (1981). Optimal Auctions. *The American Economic Review*, 71(3):381–392.
- Roughgarden, T., Talgam-Cohen, I., and Yan, Q. (2012). Supply-limiting mechanisms. In *ACM Conference on Electronic Commerce*, pages 844–861.
- Wilson, R. (1987). Game theoretic analysis of trading processes. *Advances in Economic Theory*.

A Proofs from Section 3

Lemma 3.2 (Restatement). *A mechanism with two-price allocation rule $x = x_{\text{val}} - x_C$ is BIC if and only if for all v and v^+ such that $v < v^+ \leq v + C$,*

$$\frac{x_{\text{val}}(v)}{C} \leq \frac{x_C(v^+) - x_C(v)}{v^+ - v} \leq \frac{x(v^+)}{C}. \quad (1)$$

Proof of Lemma 3.2. Consider an agent and fix two possible values of the agent $v \leq v^+ \leq v + C$. The utility for truthtelling with value v is $C \cdot x_C(v)$ in a two-price auction. The utility for misreporting v^+ from value v is $x_{\text{val}}(v^+) \cdot (v - v^+) + x_C(v^+) \cdot (C + v - v^+)$: when the mechanism sells and charges v^+ , the agent’s utility is $v - v^+$; when the mechanism sells and charges $v^+ - C$, her utility is $u_C(C + v - v^+) = C + v - v^+$ (since $v < v^+$). Likewise, the utility for misreporting v from true value v^+ is $x_{\text{val}}(v) \cdot (v^+ - v) + x_C(v) \cdot C$. Note that here when the mechanism charges $v - C$, the utility of the agent is C because the wealth $C - v + v^+$ is more than C ; when the mechanism charges v , her utility is $v^+ - v$ because we assumed $v^+ \leq v + C$.

An agent with valuation v (or v^+) would not misreport v^+ (or v) if and only if

$$x_C(v) \cdot C \geq x_{\text{val}}(v^+) \cdot (v - v^+) + x_C(v^+) \cdot (C + v - v^+); \quad (12)$$

$$x_C(v^+) \cdot C \geq x_{\text{val}}(v) \cdot (v^+ - v) + x_C(v) \cdot C. \quad (13)$$

Now the right side of (1) follows from (12) and the left side follows from (13).

When $v^+ > v + C$, the agent with value v certainly has no incentive to misreport v^+ , since any outcome results in non-positive utility. Alternatively, the agent with value v^+ will derive utility $C \cdot x(v)$ from misreporting v and thus will misreport if and only if $x(v) > x_C(v^+)$. Substituting $v + C$ for v^+ in equation (1) gives $x(v) \leq x_C(v + C)$, and taking this for intermediate points between $v + C$ and v^+ gives monotonicity of $x_C(v)$ over $[v + C, v^+]$. Combining these gives $x(v) \leq x_C(v + C) \leq x_C(v)$ and hence v^+ will not misreport v . \square

Corollary 3.3 (Restatement): *The allocation rules x_C and x_{val} of a BIC two-priced mechanism satisfies that for all $v < v^+$,*

$$\int_v^{v^+} \frac{x_{\text{val}}(z)}{C} dz \leq x_C(v^+) - x_C(v) \leq \int_v^{v^+} \frac{x(z)}{C} dz. \quad (4)$$

Proof of Corollary 3.3. Without loss of generality, suppose $v^+ \leq v + C$ (the statement then follows for higher v^+ by induction). Define function

$$\bar{x}_C(z) = x_C(v) + \int_v^z \frac{\sup_{y' \in [v, y]} x_{\text{val}}(y')}{C} dy, \quad \forall z \in [v, v^+],$$

then $\bar{x}_C(z) \geq x_C(v) + \int_v^z \frac{x_{\text{val}}(y)}{C} dy$ and hence

$$\int_v^z \frac{x_{\text{val}}(y)}{C} dy \leq \bar{x}_C(z) - x_C(v).$$

By the argument in the proof of Lemma 4.3, part (b), we have $\bar{x}_C(z) \leq x_C(z)$, for all z . This gives the left side of (4). The other side is proven similarly. \square

B Proofs from Section 4

Definition 4.1 (Restatement). We define $\bar{x} = \bar{x}_C + \bar{x}_{\text{val}}$ as follows:

- (a) $\bar{x}_{\text{val}}(v) = x_{\text{val}}(v)$;
- (b) Let $r(v)$ be $\frac{1}{C} \sup_{z \leq v} x_{\text{val}}(z)$, and let

$$\bar{x}_C(v) = \begin{cases} \int_0^v r(y) dy, & v \in [0, C]; \\ x_C(v), & v > C. \end{cases} \quad (5)$$

Lemma 4.3 (Restatement).

- (a) On $v \in [0, C]$, $\bar{x}_C(\cdot)$ is a convex, monotone increasing function.
- (b) On all v , $\bar{x}_C(v) \leq x_C(v)$.
- (c) The incentive constraint from the left-hand side of (4) holds for \bar{x}_C : $\int_v^{v^+} \bar{x}_{\text{val}}(z) dz \leq \bar{x}_C(v^+) - \bar{x}_C(v)$ for all $v < v^+$.
- (d) On all v , $\bar{x}_C(v) \leq x_C(v)$, $\bar{x}(v) \leq x(v)$, and $\bar{p}(v) \geq p(v)$.

Proof of Lemma 4.3.

- (a) On $[0, C]$, $\bar{x}_C(v)$ is the integral of a monotone, non-negative function.

- (b) The statement holds directly from the definition for $v > C$; therefore, fix $v \leq C$ in the argument below.

Since $r(v)$ is an increasing function of v , it is Riemann integrable (and not only Lebesgue integrable).

Fixing v , we show that, given any $\epsilon > 0$, $\bar{x}_C(v) \leq x_C(v) + \epsilon$. Fix an integer $N > v/\epsilon$, and let Δ be $v/N < \epsilon$. Consider Riemann sum $S = \sum_{j=1}^N \Delta \cdot r(\xi_j)$, where each ξ_j is an arbitrary point in $[(j-1)\Delta, j\Delta]$.⁵ We will also denote by $S(k) = \sum_{j=1}^k \Delta \cdot r(\xi_j)$, $k \leq N$, the partial sum of the first k terms. Since $\bar{x}_C(v) = \lim_{\Delta \rightarrow 0} S$, it suffices to show that for all $\Delta < \epsilon$, $S \leq x_C(v) + \epsilon$. In order to show this, we define a piecewise linear function y . On $[0, \Delta]$, y is 0, and then on interval $[j\Delta, (j+1)\Delta]$, y grows at a rate $r((j-1)\Delta)$. Intuitively, y “lags behind” x_C by an interval Δ and we will show it lower bounds x_C and upper bounds $S + \epsilon$. Note that since r is an increasing function, y is convex.

We first show $y(v) \leq x_C(v)$. We will show by induction on j that $y(z) \leq x_C(z)$ for all $z \in [0, j\Delta]$. Since y is 0 on $[0, \Delta]$, the base case $j = 1$ is trivial. Suppose we have shown $y(z) \leq x_C(z)$ for all $z \in [0, (j-1)\Delta]$, let us consider the interval $[(j-1)\Delta, j\Delta]$. Let z^* be $\arg \max_{z \leq (j-1)\Delta} x_{\text{val}}(z)$.⁶ By the induction hypothesis, $y(z^*) \leq x_C(z^*)$. Recall that $z^* \leq z \leq C$. By the BIC condition (2), for all $z \geq z^*$,

$$x_C(z) \geq x_C(z^*) + \frac{x_{\text{val}}(z^*)}{C}(z - z^*).$$

On the other hand, by definition, r is constant on $[z^*, z]$, and the derivative of y is no larger than $r(z^*)$ on $[z^*, z]$. Hence for all $z \leq j\Delta$,

$$\begin{aligned} y(z) &\leq y(z^*) + \frac{x_{\text{val}}(z^*)}{C}(z - z^*) \\ &\leq x_C(z^*) + \frac{x_{\text{val}}(z^*)}{C}(z - z^*) \leq x_C(z). \end{aligned}$$

This completes the induction and shows $y(z) \leq x_C(v)$ for all $z \in [0, v]$.

Now we show $S \leq y(v) + \epsilon$. Note that since $r(z) \leq 1$ for all z , $S \leq S(N-1) + \Delta < S(N-1) + \epsilon$. We will show by induction that $S(N-1) \leq y(v)$. Our induction hypothesis is $S(j-1) \leq y(j\Delta)$. The base case for $j = 1$ is obvious as $S(0) = y(\Delta) = 0$.

$$\begin{aligned} S(j) &= S(j-1) + \Delta \cdot r(\xi_j) \\ &\leq y(j) + \Delta \cdot r(j\Delta) \\ &= y(j+1). \end{aligned}$$

In the inequality we used the induction hypothesis and the monotonicity of r . The last equality is by definition of y .

This completes the proof of part (b).

⁵Obviously S depends both on Δ and the choice of ξ_j 's. For cleanness of notation we omit this dependence and do not write $S_{\Delta, \xi}$.

⁶Here we assumed that $\sup_{z \leq (j-1)\Delta} x_{\text{val}}(z)$ can be attained by z^* , which is certainly the case when x_{val} is continuous. It is straightforward to see though that we do not need such an assumption. It suffices to choose z^* such that $x_{\text{val}}(z^*)$ is close enough to $r((j-1)\Delta)$. The proof goes almost without change, except with an even smaller choice of Δ .

(c) For $v \leq v^+ \leq C$, by definition of \bar{x}_C ,

$$\bar{x}_C(v^+) - \bar{x}_C(v) = \int_v^{v^+} r(z) dz \geq \int_v^{v^+} \frac{x_{\text{val}}(z)}{C} dz.$$

For $C \leq v \leq v^+$, \bar{x}_C and \bar{x}_{val} are equal to \bar{x}_C and \bar{x}_{val} on $[v, v^+]$, and the inequality follows from Corollary 3.3. For $v \leq C$ and $v^+ \geq C$, we have

$$\begin{aligned} \bar{x}_C(v^+) - \bar{x}_C(v) &= [\bar{x}_C(v^+) - \bar{x}_C(C)] + [\bar{x}_C(C) - \bar{x}_C(v)] \\ &\geq \int_C^{v^+} \frac{x_{\text{val}}(z)}{C} dz + \int_v^C \frac{x_{\text{val}}(z)}{C} dz \\ &= \int_v^{v^+} \frac{x_{\text{val}}(z)}{C} dz. \end{aligned}$$

(d) The first part, $\bar{x}_C(v) \leq x_C(v)$, is from part (b) of the lemma and the definition of $\bar{x}_C(v) = x_C(v)$ on $v > C$. The second part, $\bar{x}(v) \leq x(v)$, follows from the definition of $\bar{x}_{\text{val}}(v) = x_{\text{val}}(v)$, the first part, and the definition of $x(v) = x_{\text{val}}(v) + x_C(v)$. The third part, $\bar{p}(v) \geq p(v)$, follows because lowering $x_C(v)$ to $\bar{x}_C(v)$ on $v \in [0, C]$ foregoes payment of $v - C$ which is non-positive (for $v \in [0, C]$). \square

C Proofs from Section 5

Theorem 5.1 (Restatement). *An allocation rule x and payment rule p are the BNE of a one-priced mechanism if and only if (a) x is monotone non-decreasing and (b) if $p(v) \geq p^{\text{VC}}(v)$ for all v then $p = p^C$ is defined as*

$$p^C(0) = 0, \tag{10}$$

$$p^C(v) = \max \left(p^{\text{VC}}(v), \sup_{v^- < v} \left\{ p^C(v^-) + (p^{\text{RN}}(v) - p^{\text{RN}}(v^-)) \right\} \right). \tag{11}$$

Moreover, if x is strictly increasing then $p(v) \geq p^{\text{VC}}(v)$ for all v and $p = p^C$ is the unique equilibrium payment rule.

The proof follows from a few basic conditions. First, with strictly monotone allocation rule x , the payment upon winning must be at least $v - C$; otherwise, a bidder would wish to overbid and see a higher chance of winning, with no decrease in utility on winning. Second, when the payment on winning is strictly greater than $v - C$, the bidder is effectively risk-neutral and the risk-neutral payment identity must hold locally. Third, when an agent is paying exactly $v - C$ on winning, they are capacitated when considering underbidding, but risk-neutral when considering overbidding. As a result, at such a point, p^C must be at least as steep as p^{RN} , i.e., if $\frac{d}{dv} p^{\text{RN}}(v) > \frac{dp^{\text{VC}}}{dv}(v)$, p^C will increase above p^{VC} , at which point it must follow the behavior of p^{RN} .

Theorem 5.1 follows from the following three lemmas which show the necessity of monotonicity, the (partial) necessity of the payment identity, and then the sufficiency of monotonicity and the payment identity.

Lemma C.1. *If x and p are the BNE of a one-priced mechanism, then x is monotone non-decreasing.*

Lemma C.1 shows that monotonicity of the allocation rule is necessary for BNE in a one-priced mechanism. Compare this to Example 3.1 where we exhibited a non-one-priced mechanisms that was not monotone. Because the utilities may be capacitated, the standard risk-neutral monotonicity argument; which involves writing the IC constraints for a high-valued agent reporting low and a low-valued agent reporting high, adding, and canceling payments; does not work.

Lemma C.2. *If x and p are the BNE of a one-priced mechanism and $p(v) \geq p^{\text{VC}}(v)$ for all v , then $p = p^C$ (as defined in Theorem 5.1); moreover, if x is strictly monotone then $p(v) \geq p^{\text{VC}}(v)$ for all v .*

From Lemma C.2 we see that one-priced mechanisms almost have a payment identity. It is obvious that a payment identity does not generically hold as a capacitated agent with value v is indifferent between payments less than $v - C$; therefore, the agent's incentives does not pin down the payment rule if the payment rule ever results in a wealth for the agent of more than C . Nonetheless, the lemma shows that this is the only thing that could lead to a multiplicity of payment rules. Additionally, the lemma shows that if x is strictly monotone, then these sorts of payment rules cannot arise.

Lemma C.3. *If allocation rule x is monotone non-decreasing and payment rule $p = p^C$ (as defined in Theorem 5.1), then they are the Bayes-Nash equilibrium of a one-priced mechanism.*

The following claim and notational definition will be used throughout the proofs below.

Claim C.4. *Compared to the wealth of type v on truthtelling, when type $v^+ > v$ misreports v she obtains strictly more wealth (and is more capacity constrained) and when type $v^- < v$ misreports v she obtains strictly less wealth (and is less capacity constrained) and if $p(v) \geq p^{\text{VC}}(v)$ then type v^- is strictly risk neutral on reporting v .*

Definition C.1. Denote the utility for type v misreporting v' for the same implicit allocation and payment rules by $U^C(v, v')$ and $U^{\text{RN}}(v, v')$ for risk-averse and risk-neutral agents, respectively.

Proof of Lemma C.1. We prove via contradiction. Assume that x is not monotone, and hence there is a pair of values, $v^- < v^+$, for which $x(v^-) > x(v^+)$. We will consider this in three cases: (1) when a type of v^- is capacitated upon truthfully reporting and winning, and when a type of v^- is strictly in the risk-neutral section of her utility upon winning and a type of v^+ is either in the (2) capacitated or (3) strictly risk-neutral section of her utility upon winning.

- (a) (v^- capacitated). If v^- is capacitated upon winning, then v^+ will also be capacitated upon winning and misreporting v^- (Claim C.4). A capacitated agent is already receiving the highest utility possible upon winning. Therefore, v^+ strictly prefers misreporting v^- as such a report (strictly) increases probability of winning and (weakly) increases utility from winning.
- (b) (v^- risk-neutral, v^+ capacitated). We split this case into two subcases depending on whether the agent with type v^- is capacitated with misreport v^+ .

- (a) (v^- capacitated when misreporting v^+). As the truthtelling v^+ type is also capacitated (by assumption of this case), the utilities of these two scenarios are the same, i.e.,

$$U^C(v^-, v^+) = U^C(v^+, v^+). \quad (14)$$

Since type v^- truthfully reporting v^- is strictly uncapacitated, if her value was increased she would feel a change in utility (for the same report); therefore, type v^+ reporting v^- has strictly more utility (Claim C.4), i.e.,

$$U^C(v^+, v^-) > U^C(v^-, v^-). \quad (15)$$

Combining (14) and (15) we arrive at the contradiction that type v^+ strictly prefers to report v^- , i.e.,

$$U^C(v^+, v^-) > U^C(v^+, v^+).$$

- (b) (v^- risk-neutral when misreporting v^+). First, it cannot be that the bidder of type v^+ is capacitated for both reports v^+ and v^- as, otherwise, misreporting v^- gives the same utility upon winning but strictly higher probability of winning. Therefore, both types are risk neutral when reporting v^- . Type v^- is risk-neutral for both reports so she feels the discount in payment from reporting v^+ instead of v^- linearly; type v^+ feels the discount less as she is capacitated at v^+ . On the other hand, v^+ has a higher value for service and therefore feels the higher service probability from reporting v^- over v^+ more than v^- . Consequently, if v^- prefers reporting v^- to v^+ , then so must v^+ (strictly).
- (c) (v^- risk-neutral, v^+ risk-neutral). First, note that the price upon winning must be higher when reporting v^- than v^+ , i.e., $p(v^-)/x(v^-) > p(v^+)/x(v^+)$; otherwise a bidder of type v^+ would always prefer to report v^- for the higher utility upon winning and higher chance of winning. Thus, a bidder of type v^+ must be risk-neutral upon underreporting v^- and winning; furthermore, risk-neutrality of v^+ for reporting v^+ implies the risk-neutrality of v^- for reporting v^+ (Claim C.4). As both v^+ and v^- are risk-neutral for reporting either of v^- or v^+ , the standard monotonicity argument for risk-neutral agents applies.

Thus, for x to be in BNE it must be monotone non-decreasing. \square

Proof of Lemma C.2. First we show that if x is strictly monotone then $p(v) \geq p^{\text{VC}}(v)$ for all v . If $p(v) < p^{\text{VC}}(v)$ then type v on truthtelling obtains a wealth w strictly larger than C . Type $v^- = v - \epsilon$, for $\epsilon \in (0, w - C)$, would also be capacitated when reporting v ; therefore, by strict monotonicity of x such a overreport strictly increases her utility and BIC is violated.

The following two claims give the necessary condition.

$$p^C(v) \geq p^C(v^-) + (p^{\text{RN}}(v) - p^{\text{RN}}(v^-)), \quad \forall v^- < v \quad (16)$$

$$p^C(v) \leq \sup_{v^- < v} \left\{ p^C(v^-) + (p^{\text{RN}}(v) - p^{\text{RN}}(v^-)) \right\}, \quad \forall v \text{ s.t. } p^C(v) > p^{\text{VC}}(v). \quad (17)$$

Equation (16) is easy to show. Since $p^C(v) \geq p^{\text{VC}}(v)$, the wealth of any type v^- when winning is at most C , and strictly smaller than C if overbidding. In other words, when overbidding, a bidder

only uses the linear part of her utility function and therefore can be seen as risk neutral. Equation (16) then follows directly from the standard argument for risk neutral agents.⁷

Equation (17) would be easy to show if p^C is continuous: for all v where $p^C(v) > p^{VC}(v)$, there is a neighborhood $(v - \epsilon, v]$ such that deviating on this interval only incurs the linear part of the utility function and the agent is effectively risk neutral. We give the following general proof that deals with discontinuity and includes continuous cases as well.

To show (17), it suffices to show that, for each v where $p^C(v) > p^{VC}(v)$, for any $\epsilon > 0$, $p^C(v) < p^C(v^-) + (p^{RN}(v) - p^{RN}(v^-) + \epsilon)$ for some $v^- < v$. Consider any $v^- > v - \frac{\epsilon}{2}$. Since $p^C(v^-) \geq p^{VC}(v^-) = (v^- - C)x(v^-) > (v - \frac{\epsilon}{2} - C)x(v^-)$, the utility for v to misreport v^- , i.e., $U^C(v, v^-)$ is not much smaller than if the agent is risk neutral:

$$U^{RN}(v, v^-) - U^C(v, v^-) < \frac{\epsilon}{2}x(v^-).$$

The following derivation, starting with the BIC condition, gives the desired bound:

$$\begin{aligned} 0 &\leq U^C(v, v) - u_C(v, v^-) < u_C(v, v) - U^{RN}(v, v^-) + \frac{\epsilon}{2}x(v^-) \\ &= (x(v)v - p^C(v)) - (x(v^-)v - p^C(v^-)) + \frac{\epsilon}{2}x(v^-) \\ &= (x(v) - x(v^-))v - (p^C(v) - p^C(v^-)) + \frac{\epsilon}{2}x(v^-) \\ &\leq p^{RN}(v) - p^{RN}(v^-) + (v - v^-)x(v) - (p^C(v) - p^C(v^-)) + \frac{\epsilon}{2}x(v^-) \\ &\leq p^{RN}(v) - p^{RN}(v^-) - (p^C(v) - p^C(v^-)) + \epsilon. \end{aligned}$$

The first equality holds because $p^C(v) > p^{VC}(v)$; the second to last inequality uses the definition of risk neutral payments (Theorem 2.1, part (b)), and the last holds because $x(v^-) < x(v) \leq 1$. \square

Proof of Lemma C.3. The proof proceeds in three steps. First, we show that an agent with value v does not want to misreport a higher value v^+ . Second, we show that the expected payment on winning, i.e., $p^C(v)/x(v)$ is monotone in v . Finally, we show that the agent with value v does not want to misreport a lower value v^- . Recall in the subsequent discussion that p^{RN} is the risk-neutral expected payment for allocation rule x (from Theorem 2.1, part (b)).

- (a) (Type v misreporting v^+ .) This argument pieces together two simple observations. First, Claim C.4 and the fact that $p^C \geq p^{VC}$ imply that v is risk-neutral upon reporting v^+ . Second, by definition of p^C , the difference in a capacitated agent's payments given by $p^C(v^+) - p^C(v)$ is at least that for a risk neutral agent given by $p^{RN}(v^+) - p^{RN}(v)$. The risk-neutral agent's utility is linear and she prefers reporting v to v^+ . As the risk-averse agent's utility is also linear for payments in the given range and because the difference in payments is only increased, then the risk-averse agent must also prefer reporting v to v^+ .
- (b) (Monotonicity of p^C/x .) The monotonicity of $\frac{p^C}{x}$, which is part (b) of Lemma 5.2, will be used in the next case (and some applications of Theorem 5.1). We consider v and v^+ and argue that $\frac{p^C(v)}{x(v)} \leq \frac{p^C(v^+)}{x(v^+)}$. First, suppose that the wealth upon winning of an agent with

⁷For a risk neutral agent, the risk neutral payment maintains the least difference in payment to prevent all types from overbidding.

value v is C , i.e., $p^C(v) = p^{\text{VC}}(v)$. If $p^C(v^+) = p^{\text{VC}}(v^+)$ as well, then by definition of p^{VC} (by $\frac{p^{\text{VC}}(v)}{x(v)} = v - C$) monotonicity of p^C/x holds for these points. If p^C is higher than p^{VC} at v^+ then this only improves p^C/x at v^+ . Second, suppose that the wealth of an agent with value v is strictly larger than C , meaning this agent's utility increases with wealth. The allocation rule $x(\cdot)$ is weakly monotone (Lemma C.1), suppose for a contradiction that $\frac{p^C(v)}{x(v)} > \frac{p^C(v^+)}{x(v^+)}$ on $v < v^+$. Then the agent with value v can pretend to have value v^+ , obtain at least the same probability of winning, and obtain strictly lower payment. This increase in wealth is strictly desired, and therefore, this agent strictly prefers misreport v^+ . Combined with part (a), above, which argued that a low valued agent would not prefer to pretend to have a higher value, this is a contradiction.

- (c) (Type v misreporting v^- .) If $p^C(v) = p^{\text{VC}}(v)$, then paying less on winning does not translate into extra utility, and hence by the monotonicity of p^C/x , the agent would never misreport.

We thus focus then on the case that $p^C(v) > p^{\text{VC}}(v)$. By the monotonicity of p^C/x , there is a point $v_0 < v$ such that for every value v^- between v_0 and v , if an agent with value v reported v^- , she would still be in the risk-neutral section of her utility function. Specifically, this entails that $\forall v^-$ such that $v_0 < v^- < v$, $p^C(v^-)/x(v^-) \geq v - C$. Consider such a v_0 and any such v^- . For any such point, $p^C(v^-)/x(v^-) > v^- - C$, and hence a bidder with value v^- would also be strictly in the risk-neutral part of her utility function upon winning.

For every such point, by our formulation in (11), $p^C(v) - p^C(v^-) = p^{\text{RN}}(v) - p^{\text{RN}}(v^-)$. As a result, since she is effectively risk-neutral in this situation, she cannot wish to misreport v^- ; otherwise, the combination of x and p^{RN} would not be BIC for risk-neutral agents.

For any $v^- \leq v_0$, the wealth on winning for a bidder with value v would increase, but only into the capacitated section of her utility function, hence gaining no utility on winning, but losing out on a chance of winning thanks to the weak monotonicity of x . Hence, she would never prefer to bid v^- over bidding v_0 . Combining this argument with the above argument, our agent with value v does not prefer to misreport any $v^- < v$. \square